Thank you for the opportunity to comment. I am choosing to remain anonymous because I am also a federal employee. Attached are some of the regulatory concerns I foresee, some suggestions, and some honest feedback on the fact that AI systems are developing - and creating unseen harms - far more quickly than our ability to regulate.

Regulatory Concerns:

1. Employment and Universal Stability Income: For all of us human workers who depend on our pay checks to pay for our human needs, ironically including software developers and coders, AI is fueling a race to the bottom. A year or two from now, only the most talented are kept on - five years from now, no one. Given the opportunity to choose between a costly human being and a free machine, companies will choose the machine that doesn't need payment or sleep. IBM has already announced it plans to pause hiring for roles it believes AI will/should fill instead. What comes in its place? How do those of us who comprise the middle and lower income classes pay for our housing, our food, our utilities, our lives? How do you prevent mass poverty and the violence that we know attends it? Some technologists have thrown their weight behind UBI. UBI is a nice idea, but dead on arrival in the U.S. context. Consider the Biden Administration's proposal to forgive 10k in student loans. It makes economic sense, we can afford it, it will make the most difference for the least privileged - and it will not go forward. UBI, moreover, isn't enough in a world with no jobs for those who don't own the robots - you'd need something closer to universal stability income to unlock the kind of human creativity "could" come with AI.

Regulations should therefore pre-emptively address mass redistribution of profits a company makes due to AI, tied directly to the employees it displaces. How you make that work in practice - and how you get something like that to pass - is impossible to prescribe, and will be cause for laughter on Capitol Hill. But AI development regulations need to include tax code adjustments, for one, support for overhauling our entire education system and requiring companies to upskill employees, moving to shorter work weeks as opposed to slashing employment. Companies (and Capitol Hill) will fight this, which is exactly the point - people need to matter. Requiring companies to share all of their profits would help companies instead point their research into finding solutions for real problems like gun violence, trafficking-in-persons, Alzheimer's and other medical issues, political polarization, and climate change.

2. Manipulation: The individual and societal implications of the torrent of disinfo that is being unleashed on the world are just beginning, and right before the 2024 election. At an individual level, this is already being used to traumatize women via the humiliation of being featured in fake porn; simulated active rapes of women, girls, and children are likely already proliferating. If the world lost its mind when the Pope "appeared" in a white puffy vest, it is hard to imagine the damage we will all suffer when video "leaks" of Putin and Xi hosting a joint meeting with their top nuclear advisors, or when QAnon supporters engage in the 2024 election. As we've seen with social media, humans are programmed to seek confirming information, not nuance - and our brains are incredibly susceptible to propaganda.

Regulations should of course mandate watermarks, hashing, and/or digital fingerprints on every AI-produced product, but this isn't enough, as Andrew Ng has pointed out. Regulations need to stipulate explicitly that Section 230 does not apply to AI-created products - and companies and platforms should be held directly responsible for the lies they propagate, the fakes they perpetuate, and the harms they create. This is the only way companies will actually police the content their technology creates - I

imagine this is one use for blockchain that could demonstrate where the original code or platform originates. This will also inflame a debate around free speech and the First Amendment- another sign that we're building technology faster than we have the capacity to adapt or stay safe.

3. Privacy and Data Use: Every person should own all of their data and be able to make decisions about how that is shared. The algorithms that sell us curated advertisements are horrifying imagined coupled with an AI tool - and that's just separating you from your money. It's not hard to imagine AI being used as a surveillance mechanism - for example, an abuser tracking a survivor of domestic violence trying to seek safety. In the hands of an authoritarian government, it's not hard to imagine AI tracking Black people, gay people, trans people, Jewish people, college students, Muslim people, immigrants, liberals - anyone who doesn't fit in their ideal society. The more quotidian worries of identify theft feel almost quaint in comparison; I'm not sure two-factor authorization or "strong passwords" hold much sway against a voice or face simulator.

When it comes to regulating and protecting privacy, the EU's GDPR is instructive. At baseline, in order to actually address privacy, we need a comprehensive federal law that covers all sectors, not a patchwork of state laws that cover only health care, data collection for those under the age of 13, or financial institutions. As a floor, data privacy needs to address: 1. data collection and the right NOT to sell or share your data with third parties; 2. data minimization - companies should only be allowed to collect the data they need to provide the service they're selling; 3. non-discrimination - in addition to protecting against discriminating based on EEO categories, companies should not be allowed to discriminate or price differentiate their services based on how much data a user is willing to give up; and 4. built-in nudges to opt-in (instead of opt-out) of data sharing/selling. Finally, it should be illegal to use AI-mined data on individuals to target them for advertising schemes. Again, given Congressional lack of action in this space, this is another instance in which our need for guardrails is seriously and detrimentally outpaced by the technology and a significant reason why we need to slow the "AI race."

4. Bias and Discrimination - As an enormous body of research indicates, we systematically design for a world that discounts the lives of women, girls, and the LGBTQ. While the U.S. has some strong anti-discrimination in employment laws, these don't actually address how bias is baked into the way we collect data, who the data illuminate, who builds the systems, or where the data come from. With AI-powered decisions, we're building a super-engine onto a world that is already rife with bias. Will AI tell us that women should receive a smaller amount of UBI, because the data tell us they chronically earn less? Will AI create additional pressure for judges to harshly sentence Black men, because AI and data science inform them that they're more likely to have appeared in court before? In essence, while it was embarrassing for Facebook to label Black men as gorillas, this kind of mistake could be life-altering and even life- ending where bias rapidly compounds other decisions.

In addition to passing a comprehensive federal data privacy law, government needs to propose guidelines for basic fairness tests for all algorithms allowed to be programmed into mass-produced AI, and industry's responsibility is to help them understand how to do so. Massive transparency around what assumptions are baked into algorithms and decision-making processes needs to be required - in plain English. It may be that we need to regulate against using AI-powered decisions in law enforcement,

in hiring, and in the legal profession in order to protect human rights and maintain the illusion that we judge people against a "jury of their peers."

5. Alignment and the Future of Humanity: Multiple AI "leaders" have now acknowledged the enormous risks associated with achieving superintelligent AI. By ChatGPT's own estimate however, the research being devoted to alignment and safety is absolutely dwarfed by the investment in the next AI tool - even though it won't provide exact estimates, the funding for alignment is in the millions, while development is in the billions. Connor Leahy estimates there are fewer than 100 people working on alignment; (ChatGPT did not put a specific number, though it tells me the number is "much smaller" than those who are working on development). A plan to rein in superhuman AI using RHLF++ seems to both overestimate the plausibility a tiny, underfunded community can get this right while underestimating our timeline - the AI is simply going to outpace us.

I'm not sure how one could plausibly regulate in this area, which is perhaps the strongest and most compelling reason to shut down or at least pause AI development. The need for alignment is so far beyond the reach of one government alone, or even our existing multilateral structure (the UN is not equipped, nimble, staffed, or resourced to address AI). How do you even begin to define human values, for example? Like Gary Marcus suggests, we need a large, international, regulatory and enforcement agency that can fund promising research and has the ability to shut down systems that are deemed unsafe. One of the questions in the comment asks about risks with self-audits; when Google fires its employees for questioning the biases built into its systems, or fires an engineer for his fear that the machines they're building are sentient - a self audit seems akin to asking the wolf, who has already decimated the chicken population, please not to eat any more. Where and how can we build a fail safe "red button" to protect ourselves?

6. Environmental Impact – I find it disorienting and delusional when AI is lauded by tech leaders as a potential panacea for solving climate change, when the systems to run AI consume so much energy. We already understand how to solve climate change; we just have not demonstrated any political will to do so.

Regulations should call AI companies' bluff. If AI is being designed to "solve pressing challenges," in addition to making it illegal to mine personal data for advertising purposes, AI companies should be regulated to disclose their environmental impact, and regulators should set standards for how AI products and services may or may not impact the environment, including but not limited to energy use, water use, pollution emitted, noise created, and material use. Enforcement measures would need to be established, including at the EPA, given AI companies would be required to self-disclose their impact.